

IN THE CLAIMS

Please amend the claims as follows:

1. (Currently Amended) A Link Packet Scheduler for each virtual lane (VL) at a given port, comprising:

a N-bit counter arranged to accumulate free credits relinquished, when a data packet is removed from a receive buffer and buffer space is reclaimed as available for data packet storage, or when a link packet is received whose Flow Control Total Bytes Sent (FCTBS) field differs from actual blocks received (ABR) at the given port;

a first comparator arranged to make comparison between accumulated free credits from the N-bit counter and a programmable credit threshold;

a second comparator arranged to make comparison between a current buffer receive utilization indicating a data storage level of the receive buffer and a programmable utilization threshold; and

a logic device arranged to track a current link state of a corresponding port, to monitor amount of receive buffer resources from the first and second comparators and to schedule for transmission of [[said]] a link packet, via a physical link.

2. (Previously Presented) The Link Packet Scheduler as claimed in claim 1, wherein said logic device comprises:

a first OR gate arranged to track whether a Link State Machine transitions into one of a LinkInitialize, a LinkArm and a LinkActive state, whether a configuration strap for enabling loopback operation changes states, and whether a per-VL Link Packet Watchdog Timer expires, and generate therefrom a logic signal;

an AND gate arranged to logically combine outputs of the first comparator and the second comparator and produce another logic signal; and

a second OR gate arranged to logically combine the logic signals from the first OR gate and the AND gate and generate an indication for transmission of the link packet for the virtual lane (VL) on the given port.

3. (Previously Presented) The Link Packet Scheduler as claimed in claim 2, wherein said

Link Packet Watchdog Timer is utilized to ensure that at minimum link packets will be independently scheduled at least once every 65,536 symbol times in accordance with InfiniBand™ specification.

4. (Original) The Link Packet Scheduler as claimed in claim 1, wherein each of said credits is defined to be 64-bytes of available receive buffer space.
5. (Previously Presented) The Link Packet Scheduler as claimed in claim 1, wherein said link packet contains at least a Flow Control Total Block Sent (FCTBS) field which is generated by a transmitter logic and is set to a total number of blocks (64 bytes) transmitted since link initialization, and a Flow Control Credit Limit (FCCL) field which is generated by receiver logic and is used to grant transmission credits to a remote transmitter.
6. (Previously Presented) The Link Packet Scheduler as claimed in claim 5, wherein said Flow Control Total Block Sent (FCTBS) field and said Flow Control Credit Limit (FCCL) field are used to guarantee that data is never lost due to lack of said receive buffer at the end of the physical link.
7. (Original) The Link Packet Scheduler as claimed in claim 1, wherein said programmable utilization threshold is set such that the receive buffer has multiple data packets pending processing.
8. (Previously Presented) A data network, comprising:
 - a host system having a host-fabric adapter;
 - at least one remote system;
 - a switch fabric which interconnects said host system via said host-fabric adapter to said at least one remote system along different physical links for data communications; and
 - one or more communication ports provided in said host-fabric adapter of said host system each port including a set of transmit and receive buffers to send and receive data packets concurrently via respective transmitter and receiver at an end of a physical link, via said switched

fabric, and a flow control mechanism utilized to prevent loss of data due to receive buffer overflow at the end of said physical link, said flow control mechanism including a N-bit counter arranged to accumulate free credits relinquished, when a data packet is removed from a receive buffer and buffer space is reclaimed as available for data packet storage, or when a link packet is received whose Flow Control Total Bytes Sent (FCTBS) field differs from actual blocks received (ABR) at a given port, a first comparator arranged to make comparison between accumulated free credits from the N-bit counter and a programmable credit threshold, a second comparator arranged to make comparison between a current buffer receive utilization indicating a data storage level of the receive buffer and a programmable utilization threshold, and a logic device arranged to track a current link state of a corresponding port, to monitor the amount of receive buffer resources from the first and second comparators and to schedule for the transmission of the link packet, via the physical link.

9. (Previously Presented) A data network, comprising:

- a host system having a host-fabric adapter;
- at least one remote system;
- a switch fabric which interconnects said host system via said host-fabric adapter to said at least one remote system along different physical links for data communications; and
 - one or more communication ports provided in said host-fabric adapter of said host system each port including a set of transmit and receive buffers to send and receive data packets concurrently via respective transmitter and receiver at an end of a physical link, via said switched fabric, and a flow control mechanism utilized to prevent loss of data due to receive buffer overflow at the end of said physical link, wherein said flow control mechanism is configured to support flow control through multiple virtual lanes VLs on a given port and to perform:
 - determining when a Link State Machine transitions into one of a LinkInitialize state, a LinkArm state and a LinkActive state, or when a configuration strap for enabling loopback operation changes states;
 - if the Link State Machine transitions into one of the LinkInitialize state, the LinkArm state and a LinkActive state, or when the configuration strap for enabling loopback operation changes states, scheduling said link packet transmission for all supported virtual lanes VLs on the

given port;

if the Link State Machine does not transition into one of the LinkInitialize state, the LinkArm state and a LinkActive state, or when the configuration strap for enabling loopback operation does not change states, determining whether a Link Packet Watchdog Timer per virtual lane VL has expired;

if the per VL Link Packet Watchdog Timer has expired, scheduling a link packet transmission for a virtual lane VL corresponding to said timer;

if the per VL Link Packet Watchdog Timer has not expired, determining whether a receive buffer utilization indicating a data storage level of the receive buffer exceeds a programmable utilization threshold;

if the receive buffer utilization exceeds the programmable utilization threshold, prohibiting the link packet transmission for that virtual lane VL;

if the receive buffer utilization does not exceed the programmable utilization threshold, determining whether free credits accumulated exceeds a programmable credit threshold;

if the free credits accumulated exceeds the programmable credit threshold, scheduling the link packet transmission for that virtual lane VL; and

if the free credits accumulated does not exceed the programmable credit threshold, prohibiting the link packet transmission for that virtual lane VL.

10. (Original) The data network as claimed in claim 9, wherein each of said credits is 64-bytes from the receive buffer, and said credits are relinquished when data packets are removed from the receive buffer and that space is reclaimed as available for packet storage, or when link packets are received whose Flow Control Total Bytes Sent (FCTBS) fields differ from actual blocks received (ABR) at the given port.

11. (Previously Presented) The data network as claimed in claim 8, wherein said flow control mechanism contains a Link Packet Scheduler per virtual lane (VL) arranged to schedule a link packet transmission in accordance with InfiniBand™ specification.

12. (Previously Presented) A data network, comprising:

a host system having a host-fabric adapter;
at least one remote system;
a switch fabric which interconnects said host system via said host-fabric adapter to said at least one remote system along different physical links for data communications; and
one or more communication ports provided in said host-fabric adapter of said host system each port including a set of transmit and receive buffers to send and receive data packets concurrently via respective transmitter and receiver at an end of a physical link, via said switched fabric, and a flow control mechanism utilized to prevent loss of data due to receive buffer overflow at the end of said physical link, wherein said flow control mechanism contains a Link Packet Scheduler per virtual lane (VL) arranged to schedule a link packet transmission for a virtual lane VL corresponding to said Link Packet Scheduler, wherein said Link Packet Scheduler comprises:
a N-bit counter arranged to accumulate free credits relinquished, when a data packet is removed from a receive buffer and buffer space is reclaimed as available for data packet storage, or when the link packet is received whose Flow Control Total Bytes Sent (FCTBS) field differs from actual blocks received (ABR) at a given port;
a first comparator arranged to make comparison between accumulated free credits from the N-bit counter and a programmable credit threshold;
a second comparator arranged to make comparison between a current buffer receive utilization indicating a data storage level of the receive buffer and a programmable utilization threshold; and
a logic device arranged to track a current link state of a corresponding port, to monitor amount of receive buffer resources from the first and second comparators and to schedule transmission of the link packet, via the physical link.

13. (Previously Presented) The data network as claimed in claim 12, wherein said logic device comprises:

a first OR gate arranged to track whether a Link State Machine transitions into one of a LinkInitialize, a LinkArm and a LinkActive state, whether a configuration strap for enabling loopback operation changes states, and whether a per-VL Link Packet Watchdog Timer expires

at a predetermined symbol time, and to produce therefrom another logic signal;

an AND gate arranged to logically combine outputs of the first comparator and the second comparator and to produce a logic signal; and

a second OR gate arranged to logically combine the logic signals from the first OR gate and the AND gate and to produce an indication for transmission of the link packet for the virtual lane (VL) on the given port.

14. (Previously Presented) The data network as claimed in claim 12, wherein a Link Packet Watchdog Timer is utilized to ensure that at minimum link packets will be independently scheduled at least once every 65,536 symbol times in accordance with InfiniBand™ specification.

15. (Original) The data network as claimed in claim 12, wherein each of said credits is defined to be 64-bytes of available receive buffer space.

16. (Previously Presented) The data network as claimed in claim 12, wherein said link packet contains at least a Flow Control Total Block Sent (FCTBS) field which is generated by a transmitter logic and is set to a total number of blocks (64 bytes) transmitted since link initialization, and a Flow Control Credit Limit (FCCL) field which is generated by a receiver logic and is used to grant transmission credits to a remote transmitter.

17. (Previously Presented) The data network as claimed in claim 16, wherein said Flow Control Total Block Sent (FCTBS) field and said Flow Control Credit Limit (FCCL) field are used to guarantee that data is never lost due to lack of said receive buffer at the end of the physical link.

18. (Original) The data network as claimed in claim 12, wherein said programmable utilization threshold is set such that the receive buffer has multiple data packets pending processing.

19. (Previously Presented) A method of flow control of a link packet in a host-fabric adapter

installed in a data network, comprising:

determining when a Link State Machine transitions into one of a LinkInitialize state, a LinkArm state and a LinkActive state, or when a configuration strap for enabling loopback operation changes state;

if the Link State Machine transitions into one of the LinkInitialize state, the LinkArm state and a LinkActive state, or when the configuration strap for enabling loopback operation changes states, scheduling transmission of a link packet for all supported virtual lanes VLs on a given port;

if the Link State Machine does not transition into one of the LinkInitialize state, the LinkArm state and a LinkActive state, or when the configuration strap for enabling loopback operation does not change states, determining whether a Link Packet Watchdog Timer per virtual lane VL has expired;

if the per VL Link Packet Watchdog Timer has expired, scheduling transmission of said link packet for a virtual lane VL corresponding to said timer;

if the per VL Link Packet Watchdog Timer has not expired, determining whether a receive buffer utilization indicating a data storage level of a corresponding receive buffer exceeds a programmable utilization threshold;

if the receive buffer utilization exceeds the programmable utilization threshold, prohibiting transmission of the link packet for that virtual lane VL;

if the receive buffer utilization does not exceed the programmable utilization threshold, determining whether free credits accumulated exceeds a programmable credit threshold;

if the free credits accumulated exceeds the programmable credit threshold, scheduling transmission of the link packet for that virtual lane VL; and

if the free credits accumulated does not exceed the programmable credit threshold, prohibiting transmission of the link packet for that virtual lane VL.

20. (Previously Presented) The method as claimed in claim 19, wherein each of said credits is 64-bytes from the receive buffer, and said credits are relinquished when data packets are removed from the receive buffer and space is reclaimed as available for packet storage, or when link packets are received whose Flow Control Total Bytes Sent (FCTBS) fields differ from actual

blocks received (ABR) at a given port.

21. (Previously Presented) The method as claimed in claim 19, wherein said Link Packet Watchdog Timer is utilized to ensure that at minimum link packets will be independently scheduled at least once every 65,536 symbol times in accordance with InfiniBand™ specification.

22. (Original) The method as claimed in claim 19, wherein each of said credits is defined to be 64-bytes of available receive buffer space.

23. (Previously Presented) The method as claimed in claim 19, wherein said link packet contains at least a Flow Control Total Block Sent (FCTBS) field which is generated by a transmitter logic and is set to a total number of blocks (64 bytes) transmitted since link initialization, and a Flow Control Credit Limit (FCCL) field which is generated by a receiver logic and is used to grant transmission credits to a remote transmitter.

24. (Previously Presented) The method as claimed in claim 19, wherein a Flow Control Total Block Sent (FCTBS) field and a Flow Control Credit Limit (FCCL) field are used to guarantee that data is never lost due to lack of said receive buffer at the end of the physical link.

25. (Original) The method as claimed in claim 19, wherein said programmable utilization threshold is set such that the receive buffer has multiple data packets pending processing.